

Fitting and goodness-of-fit test of non-truncated and truncated power-law distributions

Álvaro Corral[†] and Anna Deluca^{†*}

[†]*Centre de Recerca Matemàtica, Edifici C,
Campus Bellaterra, E-08193, Barcelona, Spain*

^{*}*Departament de Matemàtiques, Universitat
Autònoma de Barcelona, E-08193 Cerdanyola, Spain*

(Dated: December 27, 2012)

Abstract

Power-law distributions contain precious information about a large variety of processes in geoscience and elsewhere. Although there are sound theoretical grounds for these distributions, the empirical evidence in favor of power laws has been traditionally weak. Recently, Clauset et al. have proposed a systematic method to find over which range (if any) a certain distribution behaves as a power law. However, their method has been found to fail, in the sense that true (simulated) power-law tails are not recognized as such in some instances, and then the power-law hypothesis is rejected. Moreover, the method does not work well when extended to power-law distributions with an upper truncation. We explain in detail a similar but alternative procedure, valid for truncated as well as for non-truncated power-law distributions, based in maximum likelihood estimation, the Kolmogorov-Smirnov goodness-of-fit test, and Monte Carlo simulations. An overview of the main concepts as well as a recipe for their practical implementation is provided. The performance of our method is put to test on several empirical data which were previously analyzed with less systematic approaches. The databases presented here include the half-lives of the radionuclides, the seismic moment of earthquakes in the whole world and in Southern California, a proxy for the energy dissipated by tropical cyclones elsewhere, the area burned by forest fires in Italy, and the waiting times calculated over different spatial subdivisions of Southern California. We find the functioning of the method very satisfactory.

I. INTRODUCTION

Over the last decades, the importance of power-law distributions has continuously increased, not only in geoscience but elsewhere [1]. These are probability distributions defined by a probability density (for a continuous variable x) or by a probability mass function (for a discrete variable x) given by,

$$f(x) \propto \frac{1}{x^\alpha},$$

for $x \geq a$ and $a > 0$, with a normalization factor (hidden in the proportionality symbol \propto) which depends on whether x is continuous or discrete. In any case, normalization implies $\alpha > 1$. Sometimes power-law distributions are also called Pareto distributions [1, 2] (or Riemann zeta distributions in the discrete case [3]), although in other contexts the name Pareto is associated to a slightly different distribution [1]. So we stick to the clearer term power-law distribution.

These have remarkable, non-usual statistical properties, as are scale invariance and divergence of moments. The first one means that power-law functions (defined between 0 and ∞) are invariant under (properly performed) linear rescaling of axes (both x and f) and therefore have no characteristic scale, and hence cannot be used to define a prototype of the observable represented by x [4–7]. For example, no unit of distance can be defined from the gravitational field of a point mass (a power law), whereas a time unit can be defined for radioactive decay (an exponential function). However, as power-law distributions cannot be defined for all $x > 0$ but for $x \geq a > 0$ their scale invariance is not “complete” or strict.

A second uncommon property is the non-existence of finite moments; for instance, if $\alpha \leq 2$ not a single finite moment exists (no mean, no variance, etc.). This has important consequences, as the law of large numbers does not hold [8], p. 65, i.e., the mean of a sample does not converge as the size of the sample increases; alternatively, we can say that the sample mean tends to infinity, which is another way to understand the law of large numbers in this case [9], p. 393. If $2 < \alpha \leq 3$ the mean exists and is finite, but higher moments are infinite, which means for instance that the central limit theorem, in its classic formulation, does not apply (the mean of a sample is not normally distributed and has infinite standard deviation) [10]. Higher α ’s yield higher-order moments infinite, but then the situation is not so “critical”. Newman reviews other peculiar properties of power-law distributions, such as the 80|20 rule [6].

Although the normal (or Gaussian) distribution gives a non-zero probability that a human being is 10 m or 10 km tall, the definition of the probability density up to infinity is not traumatic at all, and the same happens with an exponential distribution and most “standard” distributions in probability theory. However, one already sees that the power-law distribution can be problematic, in particular for $\alpha \leq 2$, as it predicts an infinite mean, and for $2 \leq \alpha < 3$, as the variability of the sample mean is infinite. Of course, there can be variables having an infinite mean (we can easily simulate in our computer processes in which the time between events has an infinite mean), but in other cases, for physical reasons, the mean should be finite. In such situations a simple generalization is the truncation of the tail [1, 11, 12], yielding the truncated power-law distribution, defined in the same way as before by $f(x) \propto 1/x^\alpha$ but with $a \leq x \leq b$, with b finite, and with normalizing factor depending now on a and b (in some cases it is possible to have $a = 0$, see next section). Obviously, the existence of a finite upper cutoff b automatically leads to well-behaved moments, if the statistics is enough to “see” the cutoff; on the other hand, a range of scale invariance can persist, if $b \gg a$. What one finds in some practical problems is that the statistics is

not enough to decide which is the sample mean and one cannot easily conclude if a pure power law or a truncated power law is the right model for the data.

A well known example of (truncated or not) power-law distribution is the Gutenberg-Richter law for earthquake “size” [13–15]. If by size we understand radiated energy, the Gutenberg-Richter law implies that, in any seismically active region of the world, the sizes of earthquakes follow a power-law distribution, with an exponent $\alpha = 1 + 2B/3$ and B close to 1. In this case, scale invariance means that the simple question: *how big (in terms of radiated energy) are earthquakes there(?)*, has no possible answer. The non-convergence of the mean energy can easily be checked from data: catastrophic events such as the Sumatra-Andaman mega-earthquake of 2004 contribute to the mean much more than the previous recorded history [16]. Note that for the most common formulation of the Gutenberg-Richter law, in terms of the magnitude, earthquakes are not power-law distributed, but this is due to the fact that magnitude is an (increasing) exponential function of radiated energy, and therefore magnitude turns out to be exponentially distributed. In terms of magnitude, the statistical properties of earthquakes are trivial (well behaved mean, existence of a characteristic magnitude...), but we insist that this is not the case in terms of radiated energy.

Malamud [17] lists several other natural hazards following power-law distributions in some (physical) measure of size, such as landslides, rockfalls, volcanic eruptions [18, 19], and forest fires [20], and we can add rainfall [21, 22], tropical cyclones (roughly speaking, hurricanes) [23], auroras [24], tsunamis [25], etc. In some cases this broad range of responses is triggered simply by a small driving or perturbation (the slow motion of tectonic plates for earthquakes, the continuous pumping of solar radiation in hurricanes, etc.); then, this highly nonlinear relation between input and output can be labeled as crackling noise [26]. Notice that this does not apply for tsunamis, for instance, as they are not slowly driven (or at least not directly slowly driven).

Aschwanden [27] reviews disparate astrophysical phenomena which are distributed according to power laws, some of them related to geoscience: sizes of asteroids, craters in the Moon, solar flares, and energy of cosmic rays. In the field of ecology and close areas, the applicability of power-law distributions has been overviewed by White et al., mentioning also island and lake sizes [28]. Reference [11] provides bibliography for power-law and other heavy-tailed distributions in diverse disciplines, including hydrology, and Ref. [12] provides interesting geological examples.

A theoretical framework for power-law distributed sizes (and durations) of catastrophic phenomena not only in geoscience but also in condensed matter physics, astrophysics, biological evolution, neuroscience, and even the economy, is provided by the concept of self-organized crit-

icality, and summarized by the sandpile paradigm [4, 29–32]. However, although the ideas of self-organization and criticality are very reasonable in the context of most of the geosystems mentioned above [33–35], one cannot rule out other mechanisms for the emergence of power-law distributions [6, 32, 36, 37].

On the other hand, it is interesting to mention that, in addition to sizes and durations, power-law distributions have also been extensively reported in time between the occurrences of natural hazards (waiting times), as for instance in solar flares [38, 39], earthquakes [40–42], or solar wind [43]; in other cases the distributions contain a power-law part mixed with other factors [44–47]. Nevertheless, the possible relation with critical phenomena is not direct [48, 49]. The distance between events, or jumps, has received relatively less attention [50–52].

The importance of power-law distributions in geoscience is apparent; however, some of the evidence gathered in favor of this paradigm can be considered as “anecdotic” or tentative, as it is based on rather poor data analysis. A common practice is to find some (naive or not) estimation of the probability density or mass function $f(x)$ and plot $\ln f(x)$ versus $\ln x$ and look for a linear dependence between both variables. Obviously, a power-law distribution should behave in that way, but the opposite is not true: an apparent straight line in a log-log plot of $f(x)$ should not be considered a guarantee of an underlying power-law distribution, or perhaps the exponent obtained from there is clearly biased [28, 53–55]. But in order to discriminate between several competing theories or models, as well as in order to extrapolate the available statistics to the most extreme events, it is very important to properly fit power laws and to find the right power-law exponent (if any) [28].

The subject of this paper is a discussion on the most appropriate fitting, testing of the goodness-of-fit, and representation of power-law distributions, both non-truncated and truncated. A consistent and robust method will be checked on several examples in geoscience, including earthquakes, tropical cyclones, and forest fires. The procedure is in some points analogous to that of Clauset et al. [54], although there are variations in some key steps, in order to correct several drawbacks of the original method [21, 56]. The most important difference is in the criterion to select the range over which the power law holds. As the case of most interest in geoscience is that of a continuous random variable, the more involving discrete case will be postponed to a separate publication [57, 58].

II. POWER-LAW FITS AND GOODNESS-OF-FIT TESTS

A. Non-truncated and truncated power-law distributions

Let us consider a continuous power-law distribution, defined in the range $a \leq x \leq b$, where b can be finite or infinite and $a \geq 0$. The probability density of x is given by,

$$f(x) = \frac{\alpha - 1}{a^{1-\alpha} - 1/b^{\alpha-1}} \left(\frac{1}{x}\right)^\alpha,$$

the limit $b \rightarrow \infty$ provides the non-truncated power-law distribution if $\alpha > 1$ and $a > 0$, also called here pure power law; otherwise, for finite b one has the truncated power law, for which no restriction exists on α if $a > 0$, but $\alpha < 1$ if $a = 0$ (which is sometimes referred to as the power-function distribution [2]); the case $\alpha = 1$ needs a separate treatment, with

$$f(x) = \frac{1}{x \ln(b/a)}.$$

We will consider in principle that the distribution has a unique parameter, α , and that a and b are fixed and known values. Remember that, at point x , the probability density function of a random variable is defined as the probability per unit of the variable that the random variable lies in a infinitesimal interval around x , that is,

$$f(x) = \lim_{\Delta x \rightarrow 0} \frac{\text{Prob}[x \leq \text{random variable} < x + \Delta x]}{\Delta x},$$

and has to verify $f(x) \geq 0$ and $\int_{-\infty}^{\infty} f(x) dx = 1$, see for instance Ref. [59].

Equivalently, we could have characterized the distribution by its (complementary) cumulative distribution function,

$$S(x) = \text{Prob}[\text{random variable} \geq x] = \int_x^{\infty} f(x) dx.$$

For a truncated or non-truncated power law this leads to

$$S(x) = \frac{1/x^{\alpha-1} - 1/b^{\alpha-1}}{a^{1-\alpha} - 1/b^{\alpha-1}},$$

if $\alpha \neq 1$ and

$$S(x) = \frac{\ln(b/x)}{\ln(b/a)},$$

if $\alpha = 1$. Note that although $f(x)$ always has a power-law shape, $S(x)$ only has it in the non-truncated case ($b \rightarrow \infty$ and $\alpha > 1$); nevertheless, even not being a power law in the truncated case, the distribution is a power law, as it is $f(x)$ and not $S(x)$ which gives the name to the distribution.

B. Problematic fitting methods

Given a set of data, there are many methods to fit a probability distribution. Goldstein et al. [55], Bauke [53], White et al. [28], and Clauset et al. [54] check several methods based in the fitting of the estimated probability densities or cumulative distributions in the power-law case. As mentioned in the first section, $\ln f(x)$ is then a linear function of $\ln x$, both for non-truncated and truncated power laws. The same holds for $\ln S(x)$, but only in the non-truncated case. So, one can either estimate $f(x)$ from data, using some binning procedure, or estimate $S(x)$, for which no binning is necessary, and then fit a straight line by the least-squares method. As we find White et al.'s study the most complete, we summarize their results below, although those of the other authors are not very different.

For non-truncated power-law distributions, White et al. find that the results of the least-squares method using the cumulative distribution are reasonable, although the points in $S(x)$ are not independent and linear regression should yield problems in this case. We stress that this procedure only can work for non-truncated distributions (i.e., with $b \rightarrow \infty$), truncated ones yield bad results [12].

The least-squares method applied to the probability density $f(x)$ has several variations, depending on the way of estimating $f(x)$. Using linear binning one obtains a simple histogram, for which the fitting results are catastrophic [28, 53, 55, 60]. This is not unexpected, as linear binning of a heavy-tailed distribution can be considered as a very naive approach. If instead of linear binning one uses logarithmic binning the results improve (when done “correctly”), and are reasonable in some cases, but they still show some bias, high variance, and depend on the size of the bins. A fundamental point should be to avoid having empty bins, as they are disregarded in logscale, introducing an important bias.

In summary, methods of estimation of probability-distribution parameters based on least-squares fitting can have many problems, and usually the results are biased. Moreover, these methods do not take into account that the quantity to be fitted is a probability distribution (i.e., once the distributions are estimated, the method is the same as for any other kind of function). We are going to see that the method of maximum likelihood is precisely designed for dealing with probability distributions, presenting considerable advantages in front of the other methods just mentioned.

C. Maximum likelihood estimation

Let us denote a sample of the random variable x with N elements as x_1, x_2, \dots, x_N , and let us consider a probability distribution $f(x)$ parameterized by α . The likelihood function $L(\alpha)$ is defined as the joint probability density (or the joint probability mass function if the variable were discrete) evaluated at x_1, x_2, \dots, x_N in the case in which the variables were independent, i.e.,

$$L(\alpha) = \prod_{i=1}^N f(x_i).$$

Note that the sample is considered fixed, and it is the parameter α what is allowed to vary. In practice it is more convenient to work with the log-likelihood, the natural logarithm of the likelihood (dividing by N also, in our definition),

$$\ell(\alpha) = \frac{1}{N} \ln L(\alpha) = \frac{1}{N} \sum_{i=1}^N \ln f(x_i).$$

The maximum likelihood (ML) estimator of the parameter α based on the sample is just the maximum of $\ell(\alpha)$ as a function of α (which coincides with the maximum of $L(\alpha)$, obviously). For a given sample, we will denote the ML estimator as α_e (e is from empirical), but it is important to realize that the ML estimator is indeed a statistic (a quantity calculated from a random sample) and therefore can be considered as a random variable; in this case it is denoted as $\hat{\alpha}$. In a formula,

$$\alpha_e = \operatorname{argmax}_{\forall \alpha} \ell(\alpha),$$

where argmax refers to the argument of the function ℓ that make it maximum.

For the truncated or the non-truncated continuous power-law distribution we have, substituting $f(x)$ and introducing $r = a/b$, disregarding the case $a = 0$,

$$\ell(\alpha) = \ln \frac{\alpha - 1}{1 - r^{\alpha-1}} - \alpha \ln \frac{g}{a} - \ln a, \text{ if } \alpha \neq 1,$$

$$\ell(\alpha) = -\ln \ln \frac{1}{r} - \ln g, \text{ if } \alpha = 1;$$

g is the geometric mean of the data, $\ln g = N^{-1} \sum_{i=1}^N \ln x_i$, and the last term in each expression is irrelevant for the maximization of $\ell(\alpha)$. Remember that the distribution is only parameterized by α , whereas a and b (and r) are constant parameters; therefore, $\ell(\alpha)$ is not a function of a and b , but of α . In order to find the maximum one can derive with respect α and set the result equal to zero [1, 11],

$$\left. \frac{d\ell(\alpha)}{d\alpha} \right|_{\alpha=\alpha_e} = \frac{1}{\alpha_e - 1} + \frac{r^{\alpha_e-1} \ln r}{1 - r^{\alpha_e-1}} - \ln \frac{g}{a} = 0,$$

which constitutes the so-called likelihood equation for this problem. For a non-truncated distribution, $r = 0$, and it is clear that there is one and only one solution,

$$\alpha_e = 1 + \frac{1}{\ln(g/a)},$$

which corresponds to a maximum, as

$$L(\alpha) = e^{N\ell(\alpha)} = \frac{1}{a^N}(\alpha - 1)^N e^{-N\alpha \ln(g/a)},$$

has indeed a maximum (resembling a gamma probability density, see next subsection). Figure 1 illustrates the log-likelihood function and its derivate, for simulated power-law data.




FIG. 1: Log-likelihood $\ell(\alpha)$ and its derivative, for simulated non-truncated power-law data with exponent $\alpha = 1.15$ and $a = 0.001$.

In the truncated case it is not obvious that there is a solution to the likelihood equation; however, one can take advantage of the fact that the power-law distribution, for fixed a and b , can be viewed

as belonging to the regular exponential family, for which it is known that the maximum likelihood estimator exists and is unique. Indeed, in the single-parameter case, the exponential family can be written in the form,

$$f(x) = C^{-1}(\alpha)H(x)e^{\theta(\alpha) \cdot T(x)},$$

where both $\theta(\alpha)$ and $T(x)$ can be vectors, the former containing the parameter α of the family. Then, for $\theta(\alpha) = -\alpha$, $T(x) = \ln x$, and $H(x) = 1$ we obtain the (truncated or not) power-law distribution, which therefore belongs to the regular exponential family, which guarantees the existence of a unique ML solution.

In order to find the ML estimator of the exponent in the truncated case, we proceed, for practical reasons, by maximizing directly the log-likelihood $\ell(\alpha)$ (rather than by solving the likelihood equation). We will use the downhill simplex method [61], although any other simpler procedure should work, as the problem is one-dimensional. In this case, one needs to take care when the value of α gets very close to one in the maximization algorithm, and then replace $\ell(\alpha)$ by its limit at $\alpha = 1$,

$$\ell(\alpha) \rightarrow_{\alpha \rightarrow 1} -\ln \ln \frac{1}{r} - \alpha \ln \frac{g}{a} - \ln a,$$

which is in agreement with the likelihood function for a (truncated) power-law distribution with $\alpha = 1$.

An important property of ML estimators, not present in other fitting methods, is their invariance under re-parameterization. If instead of working with parameter α we use $v = h(\alpha)$, then, the ML estimator of v is “in agreement” with that of α , i.e., $\hat{v} = h(\hat{\alpha})$. Indeed,

$$\frac{d\ell}{dv} = \frac{d\ell}{d\alpha} \frac{d\alpha}{dv},$$

so, the maximum of ℓ as a function of v is attained at the point $h(\hat{\alpha})$, provided that the function h is “one-to-one”. Note that the parameters could be multidimensional as well. Reference [62] studies this invariance with much more care.

In their comparative study, White et al. conclude that maximum likelihood estimation outperforms the other methods, as always yields the lowest variance and bias of the estimator. This is not unexpected, as the ML estimator is, mathematically, the one with minimum variance among all asymptotically unbiased estimators. This property is called asymptotical efficiency [28, 53].

D. Standard deviation of the ML estimator

The interesting result of this subsection is the value of the uncertainty σ of $\hat{\alpha}$, given by [11],

$$\sigma = \frac{1}{\sqrt{N}} \left[\frac{1}{(\alpha_e - 1)^2} - \frac{r^{\alpha_e - 1} \ln^2 r}{(1 - r^{\alpha_e - 1})^2} \right]^{-1/2},$$

but we will see below that this formula is not necessary, as σ can be calculated from Monte Carlo simulations. The rest of the subsection is devoted to the particular derivation of σ for a non-truncated power-law distribution, and therefore can be skipped by readers not liking probability theory.

For the calculation of $\hat{\alpha}$ (the ML estimator of α) one needs to realize that this is indeed a statistic (a quantity calculated from a random sample) and therefore it can be considered as a random variable. Note that α denotes the true value of the parameter, which is unknown. It is more convenient to work with $\alpha - 1$ (the exponent of the cumulative distribution function); in the non-truncated case ($r = 0$ with $\alpha > 1$) we can easily derive its distribution. First let us consider the geometric mean of the sample, g , rescaled by the minimum value a ,

$$\ln \frac{g}{a} = \frac{1}{N} \sum_{i=1}^N \ln \frac{x_i}{a}.$$

As each x_i is power-law distributed (by hypothesis), a simple change of variables shows that $\ln(x_i/a)$ turns out to be exponentially distributed, with scale parameter $1/(\alpha - 1)$; then, the sum will be gamma distributed with the same scale parameter and with shape parameter given by N (this is the key property of the gamma distribution [63]). Therefore, $\ln(g/a)$ will follow the same gamma distribution but with scale parameter $N^{-1}(\alpha - 1)^{-1}$.

At this point it is useful to introduce the generalized gamma distribution [1, 2, 64], with density, for a random variable $y \geq 0$,

$$D(y) = \frac{|\delta|}{c\Gamma(\gamma/\delta)} \left(\frac{y}{c}\right)^{\gamma-1} e^{-(y/c)^\delta},$$

where $c > 0$ is the scale parameter and γ and δ are the shape parameters, which have to verify $0 < \gamma/\delta < \infty$ (so, the only restriction is that they have the same sign, although the previous references only consider $\gamma > 0$ and $\delta > 0$); the case $\delta = 1$ yields the usual gamma distribution and $\delta = \gamma = 1$ is the exponential one. Again, changing variables one can show that the inverse $z = 1/y$ of a generalized gamma variable is also a generalized gamma variable, but with transformed parameters,

$$\gamma, \delta, c \rightarrow -\gamma, -\delta, \frac{1}{c}.$$

So, $\hat{\alpha} - 1 = z = 1/\ln(g/a)$ will have a generalized gamma distribution, with parameters $-N$, -1 , and $N(\alpha - 1)$ (keeping the same order as above). Introducing the moments of this distribution [2],

$$\langle y^m \rangle = c^m \frac{\Gamma(\frac{\gamma+m}{\delta})}{\Gamma(\gamma/\delta)}$$

(valid for $m > -\gamma$ if $\gamma > 0$ and for $m < |\gamma|$ if $\gamma < 0$, and $\langle y^m \rangle$ infinite otherwise), we obtain the expected value of $\hat{\alpha} - 1$,

$$\langle \hat{\alpha} - 1 \rangle = \frac{N(\alpha - 1)}{N - 1}.$$

Note that the ML estimator, $\hat{\alpha}$, is biased, as its expected value does not coincide with the right value, α ; however, asymptotically, the right value is recovered. An unbiased estimator of α can be obtained for a small sample as $(1 - 1/N)\alpha_e + 1/N$, although this will not be of interest for us.

In the same way, the standard deviation of $\hat{\alpha} - 1$ (and of $\hat{\alpha}$) turns out to be

$$\sigma = \sqrt{\langle (\hat{\alpha} - 1)^2 \rangle - \langle \hat{\alpha} - 1 \rangle^2} = \frac{\alpha - 1}{(1 - 1/N)\sqrt{N - 2}},$$

which leads asymptotically to $(\alpha - 1)/\sqrt{N}$. In practice, we need to replace α by the estimated value α_e ; then, this is nothing else than the limit $r = 0$ ($b \rightarrow \infty$) of the general formula stated above for σ [11]. The fact that the standard deviation tends to zero asymptotically (together with the fact that the estimator is asymptotically unbiased) implies that any single estimation converges (in probability) to the true value, and therefore the estimator is said to be consistent.

E. Goodness-of-fit test

One can realize that the maximum likelihood method always yields a ML estimator for α , no matter which data one is using. In the case of power laws, as the data only enters in the likelihood function through its geometric mean, any sample with a given geometric mean yields the same value for the estimation, although the sample can come from a true power law or from any other distribution. So, no quality of the fit is guaranteed and thus, maximum likelihood estimation should be rather called minimum unlikelihood estimation. For this reason a goodness-of-fit test is necessary (although recent works do not take into account this fact [13, 28, 65]).

Following Goldstein et al. [55] and Clauset et al. [54] we use the Kolmogorov-Smirnov (KS) test, based on the calculation of the KS statistic or KS distance d_e between the theoretical probability distribution, represented by $S(x)$, and the empirical one, $S_e(x)$. The latter (which is an unbiased

estimator of the cumulative distribution) is given by the stepwise function

$$S_e(x) = n_e(x)/N,$$

where $n_e(x)$ is the number of data in the sample taking a value of the variable larger than or equal to x . The KS statistic is just the maximum difference, in absolute value, between $S(x)$ and $n_e(x)/N$, that is,

$$d_e = \max_{a \leq x \leq b} |S(x) - S_e(x)| = \max_{a \leq x \leq b} \left| \frac{1}{1 - r^{\alpha_e - 1}} \left[\left(\frac{a}{x} \right)^{\alpha_e - 1} - r^{\alpha_e - 1} \right] - \frac{n_e(x)}{N} \right|,$$

where the bars denote absolute value. Note that the theoretical cumulative distribution $S(x)$ is parameterized by the value of α obtained from ML, α_e . In practice, the difference only needs to be evaluated at the points x_i of the sample (as the routine `ksone` of Ref. [61] does). A more strict mathematical definition uses the supremum instead of the maximum, but in practice the maximum works perfectly. We illustrate the procedure in Fig. 2, with a simulation of a non-truncated power law.

Intuitively, if d_e is large the fit is bad, whereas if d_e is small the fit can be considered as good. But the relative scale of d_e is provided by its own probability distribution, through the calculation of a p -value. Under the hypothesis that the data follow indeed the theoretical distribution, with the parameter α obtained from our previous estimation (this is the null hypothesis), the p -value provides the probability that the KS statistic takes a value larger than the one obtained empirically, i.e.,

$$p = \text{Prob}[\text{KS statistic for power-law data (with } \alpha_e) \text{ is } > d_e];$$

then, bad fits will have rather small p -values.

It turns out that, in principle, the distribution of the KS statistic is known, at least asymptotically, independently of the underlying form of the distribution, so,

$$p_Q = Q(d_e \sqrt{N} + 0.12d_e + 0.11d_e/\sqrt{N}) = 2 \sum_{j=1}^{\infty} (-1)^{j-1} \exp[-2j^2(d_e \sqrt{N} + 0.12d_e + 0.11d_e/\sqrt{N})^2],$$

for which we use the routine in Ref. [61] (but note that Eq. (14.3.9) there is not right).

Nevertheless, this formula will not be accurate in our case, as we are “optimizing” the value of α using the same sample to which we apply the KS test, which yields a bias in the test, i.e., the formula would work for the true value of α , but not for one obtained by ML, which would yield in general a smaller KS statistic and too large p -values (because the fit for α_e is better than for




FIG. 2: Empirical (complementary) cumulative distribution and its corresponding fit for a simulated non-truncated power-law distribution with $\alpha = 1.15$ and $a = 0.001$. The point x at which the maximum difference between both curves takes place is marked as an illustration of the calculation of the KS statistic.

the true value α) [54, 55]. However, for this very same reason the formula can be useful to reject the goodness of a given fit, i.e., if p obtained in this way is already below 0.05, the true p will be even smaller and the fit is certainly bad. But the opposite is not true. In a formula,

$$\text{if } Q(d_e\sqrt{N} + 0.12d_e + 0.11d_e/\sqrt{N}) < 0.05, \text{ reject power law,}$$

otherwise, no decision can be taken yet. As a final comment, perhaps a more powerful test would be to use, instead of the KS statistic, the Kuiper's statistic [61], which is a refinement of the former one. We have not compared both choices, but it is stated by Clauset et al. that they yield very similar results.

F. The Clauset et al.'s recipe

Now we are in condition to explain the genuine Clauset et al.'s method [54]. As we are not going to apply it, readers not interested or not curious can skip the subsection. The key to fitting

a power law is neither the ML estimation of the exponent nor the goodness-of-fit test, but the selection of the interval $[a, b]$ over which the power law holds. Initially, we have taken a and b as fixed parameters, but in practice this is not the case, and one has to decide where the power law starts and where ends, independently of the total range of the data. In any case, N will be the number of data in the power-law range (and not the total number of data).

The recipe of Clauset et al. applies to non-truncated power-law distributions ($b \rightarrow \infty$), and considers that a is a variable which needs to be fit from the sample (values of x below a are outside the power-law range). The recipe simply consists in the search of the value of a which yields a minimum of the KS statistic, using as a parameter of the theoretical distribution the one obtained by maximum likelihood, α_e , for the corresponding a (no calculation of a p -value is required for each fixed a). In other words,

$$a = \text{the one that yields minimum } d_e$$

Then, a global p -value is computed by generating synthetic samples by a mixture of parametric and non-parametric bootstrap, to which the same procedure is applied in order to fit a and α (i.e., minimization of the KS distance using ML for fitting).

These authors do not provide any explanation of why this should work, although one can argue that, if the data is indeed a power law with the desired exponent, the larger the number of data (the smaller the a -value), the smaller the value of d_e , as d_e goes as $1/\sqrt{N}$ (for large N , see previous subsection). On the other hand, if for a smaller a the data departs from the power law, this deviation should compensate and overcome the reduction in d_e due to the increase of N , yielding a larger d_e . But there is no reason to justify this overcoming.

Nevertheless, we will not use the Clauset's et al.'s procedure for two other reasons. First, its extension to truncated power laws, although obvious, and justifiable with the same arguments, yields bad results, as the resulting values of the upper truncation cutoff, b , are highly unstable. Second, even for non-truncated distributions, it has been shown that the method fails to detect the existence of a power law for data simulated with a power-law tail [56]: the method yields an a -value well below the true power-law region, and therefore the resulting p is too small for the power law to become acceptable. We will explain an alternative method that avoids these problems, but first let us come back to the case with a and b fixed.

G. Monte Carlo simulations

Remember that we are considering a power law distribution, defined in $a \leq x \leq b$. We already have fit the distribution, by ML, and we are testing the goodness of the fit by means of the KS statistic. In order to obtain a reliable p -value for this test we will perform Monte Carlo simulations of the whole process. A synthetic sample power-law distributed and with N elements can be obtained in a straightforward way, from the inversion or transformation method [59, 61, 66],

$$x_i = \frac{a}{[1 - (1 - r^{\alpha_e - 1})u_i]^{1/(\alpha_e - 1)}},$$

where u_i represents a uniform random number in $[0, 1)$. Use your favorite random number generator for it [61].

H. Application of the complete procedure to the syntetic sample and calculation of p -value

The previous fitting and testing procedure is applied in exactly the same way to the synthetic sample, yielding a ML exponent α_s (from synthetic, or simulated) and a KS statistic d_s , computed as the difference between the theoretical cumulative distribution, with parameter α_s , and the simulated one, $n_s(x)/N$ (obtained from simulations with α_e , as described in the previous paragraph), i.e.,

$$d_s = \max_{a \leq x \leq b} \left| \frac{1}{1 - r^{\alpha_s - 1}} \left[\left(\frac{a}{x} \right)^{\alpha_s - 1} - r^{\alpha_s - 1} \right] - \frac{n_s(x)}{N} \right|.$$

Both values of the exponent, α_e and α_s , should be close, but not necessarily the same. Note that we are not parameterizing $S(x)$ by the empirical value α_e , but with a new fitted value α_s . This is in order to avoid biases, as a parameterization with α_e would lead to worse fits (as the best one would be with α_s) and therefore to larger values of the resulting KS statistic and to artificially smaller p -values. So, although the null hypothesis of the test is that the exponent of the power law is α_e , and synthetic samples are obtained with this value, no further knowledge of this value is used in the test. This is the procedure used in Refs. [54, 67], but it is not clear if it is the one of Ref. [55].

In fact, one single synthetic sample is not enough to do a proper comparison with the empirical sample, and we repeat the simulation many times (avoiding of course to use the same seed of the random number generator for each sample). The most important outcome is the set of values of the KS statistic, d_s , which allows to estimate its distribution. The p -value is simply calculated as

$$p = \frac{\text{number of simulations with } d_s \geq d_e}{N_s},$$

where N_s is the number of simulations. Figure 3 shows an example of the distribution of the KS statistic for simulated data, which can be used as a table of critical values when the number of data and the exponent are the same as in the example [55].




FIG. 3: Cumulative (complementary) distribution of the Kolmogorov-Smirnov statistic for simulated non-truncated power-law distributions with $\alpha = 1.15$.

The standard deviation of the p -value can be calculated just using that the number of simulations with $d_s \geq d_e$ is binomially distributed, with standard deviation $\sqrt{N_s p(1-p)}$ and therefore the standard deviation of p is the same divided by N_s ,

$$\sigma_p = \sqrt{\frac{p(1-p)}{N_s}}.$$

In fact, the p -value in this formula should be the ideal one (the one of the whole population) but we need to replace it by the estimated value; further, when doing estimation from data, N_s should be $N_s - 1$, but we have disregarded this bias correction. It will be also useful to consider the relative uncertainty of p , which is the same as the relative uncertainty of the number of simulations with $d_s \geq d_e$ (as both are proportional). Dividing the standard deviation of p by its mean (which is p),

we obtain

$$CV_p = \sqrt{\frac{1-p}{pN_s}} \simeq \sqrt{\frac{1-p}{\text{number of simulations with } d_s \geq d_e}}$$

(we will recover this formula for the error of the estimation of the probability density).

In this way, small p -values are associated to large values of d_e , and therefore to bad fits. However, note that if we put the threshold of rejection in, let us say, $p \leq 0.05$, even true power-law distributed data, with exponent α_e , yield “bad fits” in one out of 20 samples (on average). So we are rejecting true power laws in 5 % of the cases (type I error). On the other hand, lowering the threshold of rejection would reduce this problem, but would increase the probability of accepting false power laws (type II error). In this type of tests a compromise between both types of errors is always necessary, and depends on the relative costs of rejecting a true hypothesis or accepting a false one.

In addition, we can obtain from the Monte Carlo simulations the uncertainty of the ML estimator, just computing $\bar{\alpha}_s$, the average value of α_s , and from here its standard deviation,

$$\sigma = \sqrt{(\alpha_s - \bar{\alpha}_s)^2}.$$

This procedure yields good agreement with the analytical formula of Ref. [11], but can be much more useful in the discrete power-law case.

I. Alternative method to the one by Clauset et al.

At this point, for given values of the truncation points, a and b , we are able to obtain the corresponding ML estimation of the power-law exponent as well as the goodness of the fit, by means of the p -value. Now we face the same problem Clauset et al. tried to solve: how to select the values of a and b ? We adopt the simple method proposed by Peters et al. [21]: sweeping many different values of a and b we should find, if the null hypothesis is true (i.e., if the sample is power-law distributed), that many sets of intervals yield acceptable fits (high enough p -values), so we need to find the “best” of such intervals. And which one is the best? For a non-truncated power law the answer is easy, we select the largest interval, i.e., the one with the smaller a . All the other acceptable intervals will be inside this one.

But if the power law is truncated the situation is not so clear, as there can be several non-overlapping intervals. In fact, many true truncated power laws can be contained in the data, at least there are well know examples of stochastic processes with double power-law distributions

[41, 68–70]. At this point any selection can be reasonable, but if one insists in having an automatic, blind procedure, a possibility is to select either the interval which contains the larger number of data, N [21], or the one which has the larger log-range, b/a . For double power-law distributions, in which the exponent for small x is smaller than the one for large x , the former recipe has a tendency to select the first (small x) power-law regime, whereas the second procedure usually leads to the other (large x) power law.

In summary, the final step of the method for truncated power-law distributions is contained in the formula

$$[a, b] = \text{the one that yields higher } \left\{ \begin{array}{c} N \\ \text{or} \\ b/a \end{array} \right\} \text{ provided that } p > p_c,$$

which contains in fact two procedures, one maximizing N and the other maximizing b/a . We will test both in this paper. For non-truncated power-law distributions the two procedures are equivalent.

One might be tempted to choose $p_c = 0.05$, however, it is safer to consider a larger value, as for instance $p_c = 0.20$. Note that the p -value we are using is the one for fixed a and b , and then the p -value of the whole procedure should be different, but at this point it is not necessary to obtain such a p -value, as we should have already come out with a reasonable fit. Figure 4 shows the results of the method for true power-law data.

J. Truncated or non-truncated power-law distribution?

The most natural, simplest, and interesting choice is to try to fit first a non-truncated power law distribution. If an acceptable fit is found, it is expected that a truncated power law, with $b \geq x_{max}$, the largest value of x , would yield also a good fit. In fact, if b is not considered as a fixed value but as a parameter to fit, its maximum likelihood estimator when the number of data is fixed, i.e., when b is in the range $b \geq x_{max}$, is $b_e = x_{max}$. This is easy to see [11], just looking at the equation for $\ell(\alpha)$, which increases as b approaches x_{max} . (In the same way, the ML estimator of a , for fixed number of data, would be $a_e = x_{min}$, but we are not interested in such a case.) On the other hand, it is reasonable that a truncated power law yields a better fit than a non-truncated one, as the former has two parameters and the latter only one (assuming that a is fixed, in any case).




FIG. 4: Evolution as a function of a of the KS statistic, the false p -value p_Q , the true p -value, and the estimated exponent. The true exponent, here called α_t and equal to 1.15, together with a 2σ interval, are displayed as thin black lines.

In order to do a proper comparison, in such situations the so-called Akaike information criterion (AIC) can be used. This is defined simply as the difference between twice the number of parameters and twice the maximum of the log-likelihood, i.e.,

$$AIC = 2 \times (\text{number of parameters}) - 2\ell(\alpha_e).$$

In general, more parameters lead to better fits, and to higher likelihood, so, the first term compensates this factor. Therefore, given two models, the one with smaller AIC is preferred. However, for the data analyzed in this paper, the power-law regimes are usually not comparable, as they are based on different data. So, the AIC will be of little use, except if the data are exactly the same.

III. ESTIMATION OF PROBABILITY DENSITIES AND CUMULATIVE DISTRIBUTION FUNCTIONS

The method of maximum likelihood does not rely on the estimation of the probability distributions, in contrast to other methods. Nevertheless, in order to present the results, it is useful to display some representation of the distribution, together with its fit. This procedure has no statistical value (it cannot provide a substitution of a goodness-of-fit test) but is very helpful as a visual guide, specially in order to detect bugs in the algorithms.

A. Estimation of the probability density

Remember that the definition of the probability density is

$$f(x) = \lim_{\Delta x \rightarrow 0} \frac{\text{Prob}[x \leq \text{random variable} < x + \Delta x]}{\Delta x},$$

but in practice the width Δx cannot tend to zero (there would be no statistics in such case), and one has to take a non-zero value of the width. The most usual procedure is to draw a histogram using linear binning (bins of constant width); however, there is no reason why the width of the distribution should be fixed (even, some authors seem to believe that the most correct procedure is to take $\Delta x = 1$). In fact, Δx should be chosen in order to balance the necessity of having enough statistics (large Δx) with that of having a good sampling of the function (small Δx). For power-law distributions and other fat-tailed distributions, which take values across many different scales, the right choice depends of the scale of x . In this cases it is very convenient to use the so-called logarithmic binning [31, 71]. This uses bins that appear as constant in logarithmic scale, but that in fact grow exponentially (for which the method is sometimes called exponential binning instead). Curiously, this useful method is not considered by classic texts on density estimation [72].

Let us consider the semi-open intervals $[a_0, b_0), [a_1, b_1), \dots, [a_k, b_k), \dots$, also called bins, with $a_{k+1} = b_k$ and $b_k = Ba_k$ (this constant B has nothing to do with the one in the Gutenberg-Richter law, Sec. 1). For instance, if $B = \sqrt[5]{10}$ this yields 5 intervals for each order of magnitude. Notice that the width of every bin grows linearly with a_k , but exponentially with k , as $b_k - a_k = (B - 1)a_k = a_0(B - 1)B^k$. The value of B should be chosen in order to avoid low populated bins, otherwise, a spurious exponent equal to one appears [31].

We simply will count the number of occurrences of the variable in each bin. For each value of

the random variable x_i , the corresponding bin is found as

$$k = \frac{\ln(x_i/a_0)}{\ln B}.$$

Of course, a_0 has to be smaller than any possible value of x . For a continuous variable the concrete value of a_0 should be irrelevant (if it is small enough), but in practice one has to avoid that the values of a_k coincide with round values of the variable [56].

So, with this logarithmic binning, the probability density can be estimated (following its definition) as the relative frequency of occurrences in a given bin divided by its width, i.e.,

$$f_e(x_k^*) = \frac{\text{number of occurrences in bin } k}{(b_k - a_k) \times \text{number of occurrences}},$$

where the estimation of the density is associated to a value of x represented by x_k^* . The most practical solution is to take it in the middle of the interval in logscale, so $x_k^* = \sqrt{a_k b_k}$. However, for sparse data covering many orders of magnitude it is necessary to be more careful. In fact, what we are looking for is the point x_k^* whose value of the density coincides with the probability of being between a_k and b_k divided by the width of the interval. This is the solution of

$$f(x_k^*) = \frac{1}{b_k - a_k} \int_{a_k}^{b_k} f(x) dx = \frac{S(a_k) - S(b_k)}{b_k - a_k},$$

where f and S are the theoretical distributions. When the distribution and its parameters are known, the equation can be solved either analytically or numerically. It is easy to see that for a power-law distribution (truncated or not) the solution can be written

$$x_k^* = \sqrt{a_k b_k} \left[(\alpha - 1) \frac{B^{\alpha/2-1} (B - 1)}{B^{\alpha-1} - 1} \right]^{1/\alpha},$$

where we have used that $B = b_k/a_k$ (if we were not using logarithmic binning we would have to write a bin-dependent B_k). Note that for constant (bin-independent) B , i.e., for logarithmic binning, the solution is proportional but not equal to the geometric mean of the extremes of the bin. Nevertheless, the omission of the proportionality factor does not alter the power-law behavior, just shifts (in logarithmic scale) the curve. But for a different binning procedure this is no longer true. Moreover, for usual values of B the factor is very close to one [71], although large values of B [56] yield noticeable deviations if the factor in brackets is not included, see also our treatment of the radionuclide half-lives in Sec. III, with $B = 1$. Once the value of B is fixed (usually in this paper to $\sqrt[5]{10}$), in order to avoid empty bins we merge consecutive bins until the resulting merged

bins are not empty. This leads to a change in the effective value of B for merged bins, but the method is still perfectly valid.

The uncertainty of $f_e(x)$ can be obtained from its standard deviation (the standard deviation of the estimation of the density, f_e , not of the original random variable x). Indeed, assuming independence in the sample (which is already implicit in order to apply maximum likelihood estimation), the number of occurrences of the variable in bin k is a binomial random variable (in the same way as for the p -value). As the number of occurrences is proportional to $f_e(x)$, the ratio between the standard deviation and the mean for the number of occurrences will be the same as for $f_e(x)$, which is,

$$\frac{\sigma_f(x)}{f_e(x)} = \sqrt{\frac{q}{\text{mean number of occurrences in } k}} \simeq \frac{1}{\sqrt{\text{occurrences in } k}},$$

where we replace the mean number of occurrences in bin k (not available from a finite sample) by the actual value, and q , the probability that the occurrences are not in bin k , by one. This estimation of $\sigma_f(x)$ fails when the number of counts in the bin is too low, in particular if it is one.

One final consideration is that the fitted distributions are normalized between a and b , with N number of data, whereas the empirical distributions include all data, with N_{tot} of them, $N_{tot} \geq N$. Therefore, in order to compare the fits with the empirical distributions, we will plot $Nf(x)/N_{tot}$ together with $f_e(x_k^*)$.

B. Estimation of the cumulative distribution

The estimation of the (complementary) cumulative distribution is much simpler, as bins are not involved. One just needs to sort the data, in ascending order, $x_{(1)} \leq x_{(2)} \leq \dots \leq x_{(N_{tot}-1)} \leq x_{(N_{tot})}$; then, the estimated cumulative distribution is

$$S_e(x_{(i)}) = \frac{n_e(x_{(i)})}{N_{tot}} = \frac{N_{tot} - i + 1}{N_{tot}},$$

for the data points, constant below these data points, and $S_e(x) = 0$ for $x > x_{(N_{tot})}$; $n_e(x_{(i)})$ is the number of data with $x \geq x_{(i)}$ in the empirical sample. We use the case of empirical data as an example, but it is of course the same for simulated data. For the comparison of the empirical distribution with the theoretical fit we need to correct the different number of data in both cases. So, we plot both $[NS(x) + n_e(b)]/N_{tot}$ and $S_e(x)$, in order to check the accuracy of the fit.

IV. DATA ANALYZED AND RESULTS

We have explained how, in order to try to certify that a set of data is compatible with a simple power-law distribution, a large number of mathematical formulas and immense calculations are required. Now we check the performance of our method with diverse geophysical data, which were previously analyzed with different, less rigorous or worse-functioning methods. For the peculiarities and challenges of the data set, we also include the half-lives of unstable nuclides. The parameters of the method are fixed to $N_s = 1000$ Monte Carlo simulations and the values of a and b are found sweeping a fixed number of points per order of magnitude, equally spaced in logarithmic scale. This number is 10 for non-truncated power laws (in which b is fixed to infinity) and 5 for truncated power laws. Three values of p_c are considered: 0.1, 0.2, and 0.5, in order to compare the dependence of the results on this parameter.

A. Half-lives of the elements

Reference [56] studied the statistics of the half-lives of radioactive elements. The basic entity of study was the nuclide, which is defined as an isotope of a chemical element in a concrete energy state (either the fundamental one or an excited one). Nuclides are therefore specified by the number of protons, the number of neutrons, and the energy state. Roughly speaking, the term nuclide can refer both to the atomic species or to its nucleus, but the distinction will be irrelevant to us. Unstable nuclides are called radionuclides.

Any radionuclide has a constant probability of disintegration per unit time, the decay constant, let us call it λ [73]. If M is the total amount of radioactive material at time t , this means that

$$-\frac{1}{M} \frac{dM}{dt} = \lambda.$$

This leads to an exponential decay, for which a half-life $t_{1/2}$ or a lifetime τ can be defined, as

$$t_{1/2} = \tau \ln 2 = \frac{\ln 2}{\lambda}.$$

It is well known that the half-lives take disparate values, for example, that of ^{238}U is 4.47 (American) billions of years, whereas for other nuclides it is a very tiny fraction of a second.

It has been recently claimed that these half-lives are power-law distributed [56]. In fact, 3

power-law regions were identified in the probability density of $t_{1/2}$, roughly,

$$f(t_{1/2}) \propto \begin{cases} 1/t_{1/2}^{0.65} & \text{for } 10^{-6}s \leq t_{1/2} \leq 0.1s \\ 1/t_{1/2}^{1.19} & \text{for } 100s \leq t_{1/2} \leq 10^{10}s \\ 1/t_{1/2}^{1.09} & \text{for } t_{1/2} \geq 10^8s. \end{cases}$$

Notice that there is some overlap between two of the intervals, as reported in the original reference, due to problems in delimiting the transition region. The study used variations of the Clauset et al.'s method of minimization of the KS statistic [54], introducing an upper cutoff and additional conditions to escape from the global minimum of the KS statistic, which yielded the rejection ($p = 0.000$) of the power-law hypothesis. These additional conditions were of the type of taking either a or b/a greater than a fixed amount.

For comparison, we will apply instead the method explained in the previous section to this problem. Obviously, our random variable will be $x = t_{1/2}$. The data is exactly the same as in the original reference, coming from the Lund/LBNL Nuclear Data Search web page [74]. Elements whose half-life is only either bounded from below or from above are discarded for the study, which leads to 3002 radionuclides with well-defined half-lives; 2279 of them are in their ground state and the remaining 723 in an excited state. The minimum and maximum half-lives in the data set are 3×10^{-22} s and 7×10^{31} s, respectively, yielding more than 53 orders of magnitude of logarithmic range. Further details are in Ref. [56].

The results of our fitting and testing method are shown in Table I and in Fig. 5. The fitting of a non-truncated power law yields results in agreement with Ref. [56], with $\alpha = 1.09 \pm 0.01$ and $a = 3 \times 10^7$ s, for the 3 values of p_c analyzed (0.1, 0.2, and 0.5). When fitting a truncated power law, the maximization of the log-range, b/a , yields essentially the same results as for a non-truncated power law, with slightly smaller exponents α due to the finiteness of b (results not shown). In contrast, the maximization of the number of data N yields an exponent $\alpha \simeq 0.95$ between $a \simeq 0.1$ s and $b \simeq 400$ s (with some variations depending on p_c). This result is in disagreement with Ref. [56], which yielded a smaller exponent for smaller values of a and b . In fact, as the intervals do not overlap both results are compatible, but it is also likely that a different function would lead to a better fit; for instance, a lognormal between 0.01 s and 10^5 s was proposed in Ref. [56], although the fitting procedure there was not totally reliable. Finally, the intermediate power-law range reported in the original paper (the one with $\alpha = 1.19$) is not found by any of our algorithms working on the entire dataset. It is necessary to cut the data set, removing data below, for instance, 100 s

TABLE I: Results of the fits for the $N_{tot} = 3002$ nuclide half-lives data, for different values of p_c . We show the cases of a pure or non-truncated power law (with $b = \infty$, fixed) and truncated power law (with b finite, estimated from data), maximizing N .

N	a (s)	b (s)	b/a	$\alpha \pm \sigma$	p_c
143	0.316E+08	∞	∞	1.089 ± 0.007	0.10
143	0.316E+08	∞	∞	1.089 ± 0.007	0.20
143	0.316E+08	∞	∞	1.089 ± 0.008	0.50
1596	0.0794	501	6310	0.952 ± 0.010	0.10
1539	0.1259	501	3981	0.959 ± 0.011	0.20
1458	0.1259	316	2512	0.950 ± 0.011	0.50

(which is equivalent to impose $a > 100$ s), in order that the algorithm converges to that solution. So, caution must be taken when applying the algorithm blindly, as important power-law regimes may be hidden by others having either larger N or larger log-range.

B. Seismic moment of earthquakes

The statistics of the sizes of earthquakes [75] has been investigated not only since the introduction of the first magnitude scale, by C. Richter, but even before, in the 1930's, by Wadati [15]. From a modern perspective, the most reliable measure of earthquake size is given by the (scalar) seismic moment M , which is the product of mean final slip, rupture area, and the rigidity of the fault material [76]. It is usually assumed that the energy radiated by an earthquake is proportional to the seismic moment [14], so, a power-law distribution of the seismic moment implies a power-law distribution of energies, with the same exponent.

The most relevant results for the distribution of seismic moment are those of Kagan for world-wide seismicity [13], who showed that its probability density has a power-law body, with a universal exponent in agreement with $\alpha = 1.63 \simeq 5/3$, but with an extra, non-universal exponential decay (at least in terms of the complementary cumulative distribution). However, Kagan analysis, ending in 2000, refers to a period of global seismic “quiescence”; in particular, the large Sumatra-Andaman earthquake of 2004 and the subsequent global increase of seismic activity are not included. Much recently, Main et al. have shown, using a Bayesian information criterion, that the inclusion of the new events leads to the preference of the non-truncated power-law distribution




FIG. 5: Estimated probability density of the half-lives of the radionuclides, together with the power-law fits explained in the text. The number of log-bins per order of magnitude is one, which poses a challenge in the correct estimation of the exponent, as explained in Sec. II. Data below 10^{-10} s are not shown.

in front of models with a faster large- M decay.

We take the Centroid Moment Tensor (CMT) worldwide catalog analyzed by Kagan and by Main et al., including now data from January 1977 to December 2010, and apply our statistical method to it. Although the statistical analysis of Kagan is rather complete, his procedure is different to ours. Note also that the data set does not comprise the recent (2011) Tohoku earthquake in Japan, nevertheless, the qualitative change in the data with respect to Kagan's period of analysis is very remarkable. Following this author, we separate the events by their depth: shallow for depth ≤ 70 km, intermediate for $70 \text{ km} < \text{depth} \leq 300$ km, and deep for depth > 300 km. The number of earthquakes in each category is 26824, 5281, and 1659, respectively.

Second, we also consider the Southern California's Waveform Relocated Earthquake Catalog, from January 1st, 1981 to June 30th, 2011, covering a rectangular box of coordinates $(122^\circ\text{W}, 30^\circ\text{N})$, $(113^\circ\text{W}, 37.5^\circ\text{N})$ [77, 78]. This catalog contains 111981 events with $m \geq 2$. As, in contrast with the CMT catalog, this one does not report the seismic moment M , the magnitudes m

there are converted into seismic moments, using the formula

$$\log_{10} M = \frac{3}{2}(m + 6.07),$$

where M comes in units of Nm (Newtons times meters); however, this formula is a very rough estimation of seismic moment, as it is only accurate (and exact) when m is the so-called moment magnitude [14], whereas the magnitudes recorded in the catalog are of a different type. In any case, our procedure here is equivalent to fit an exponential distribution to the magnitudes reported in the catalog.

Tables II and III and Fig. 6 summarize the results of analyzing these data with our method, taking $x = M$ as the random variable. Starting with the non-truncated power-law distribution, we always obtain an acceptable (in the sense of non-rejectable) power-law fit, valid for several orders of magnitude. In all cases the exponent α is between 1.61 and 1.71, but for Southern California it is always very close to 1.66. For the worldwide CMT data the largest value of a is 3×10^{18} Nm, corresponding to a magnitude $m = 6.25$ (for shallow depth), and the smallest is $a = 8 \times 10^{16}$ Nm, corresponding to $m = 5.2$ (intermediate depth). If the events are not separated in terms of their depth (not shown), the results are dominated by the shallow case, except for $p_c = 0.5$, which leads to very large values of a and α ($a = 5 \times 10^{20}$ Nm and $\alpha \simeq 2$). The reason is probably the mixture of the different populations, in terms of depth, which is not recommended by Kagan. This is an indication that the inclusion of an upper limit b to the power law may be appropriate, with each depth corresponding to different b 's. For Southern California, the largest a found (for $p_c = 0.5$) is 1.6×10^{15} Nm, giving $m = 4$. This value is somewhat higher, in comparison with the completeness magnitude of the catalog; perhaps the reason that the power-law fit is rejected for smaller magnitudes is due to the fact that these magnitudes are not true moment magnitudes, but come from a mixture of different magnitude definitions. If the value of a is increased, the number of data N is decreased and the power-law hypothesis is more difficult to reject, due simply to poorer statistics. When a truncated power law is fitted, using the method of maximizing the number of data leads to similar values of the exponents, although the range of the fit is in some cases moved to smaller values (smaller a , and b smaller than the maximum M on the dataset). The method of maximizing b/a leads to results that are very close to the non-truncated power law.

TABLE II: Results of the non-truncated power-law fit ($b = \infty$) applied to the seismic moment of earthquakes in CMT worldwide catalog (separating by depth) and to the Southern California catalog, for different p_c .

Catalog	N	a	$\alpha \pm \sigma$	p_c
CMT deep	1216	0.1259E+18	1.622 ± 0.019	0.10
intermediate	3701	0.7943E+17	1.654 ± 0.011	0.10
shallow	5799	0.5012E+18	1.681 ± 0.009	0.10
CMT deep	898	0.1995E+18	1.608 ± 0.020	0.20
intermediate	3701	0.7943E+17	1.654 ± 0.011	0.20
shallow	5799	0.5012E+18	1.681 ± 0.009	0.20
CMT deep	898	0.1995E+18	1.608 ± 0.021	0.50
intermediate	3701	0.7943E+17	1.654 ± 0.011	0.50
shallow	1689	0.3162E+19	1.706 ± 0.018	0.50
S. California	1327	0.1000E+16	1.660 ± 0.018	10
S. California	1327	0.1000E+16	1.660 ± 0.018	20
S. California	972	0.1585E+16	1.654 ± 0.021	50

C. Energy of tropical cyclones

Tropical cyclones are devastating atmospheric-oceanic phenomena comprising tropical depressions, tropical storms, and hurricanes or typhoons [79]. Although the counts of events every year have been monitored for a long time, and other measurements to evaluate annual activity have been introduced (see Ref. [80] for an overview), little attention has been paid to the statistics of individual tropical cyclones.

In 2005, Emanuel introduced the power-dissipation index (PDI) as a simple way to obtain a rough estimation of the total energy dissipated by all tropical cyclones in a given season and some ocean basin [81]. But the PDI can also be used to characterize individual events as well, as it was done later in Ref. [23]. Indeed, the PDI is defined as the sum for all the discrete times t (that comprise the lifetime of a tropical cyclone) of the cube of the maximum sustained wind speed multiplied by the time interval of sampling, Δt . In a formula,

$$PDI = \sum_{\forall t} v_t^3 \Delta t,$$

where v_t is the maximum sustained wind speed. In the so-called best-track records, $\Delta t = 6$ hours;

TABLE III: Same as the previous table, but fitting a truncated power-law distribution, by maximizing the number of data.

Catalog	N	a	b	b/a	$\alpha \pm \sigma$	p_c
CMT deep	1216	0.1259E+18	0.3162E+23	0.2512E+06	1.621 ± 0.019	0.10
intermediate	3701	0.7943E+17	0.7943E+22	0.1000E+06	1.655 ± 0.011	0.10
shallow	13740	0.1259E+18	0.5012E+20	0.3981E+03	1.642 ± 0.007	0.10
CMT deep	1076	0.1259E+18	0.5012E+19	0.3981E+02	1.674 ± 0.033	0.20
intermediate	3701	0.7943E+17	0.7943E+22	0.1000E+06	1.655 ± 0.011	0.20
shallow	13518	0.1259E+18	0.1995E+20	0.1585E+03	1.636 ± 0.008	0.20
CMT deep	898	0.1995E+18	0.3162E+23	0.1585E+06	1.604 ± 0.021	0.50
intermediate	3701	0.7943E+17	0.7943E+22	0.1000E+06	1.655 ± 0.011	0.50
shallow	11727	0.1259E+18	0.1995E+19	0.1585E+02	1.608 ± 0.012	0.50
S. California	1146	0.1259E+16	0.1259E+22	0.1000E+07	1.663 ± 0.020	10
S. California	1146	0.1259E+16	0.1259E+22	0.1000E+07	1.663 ± 0.020	20
S. California	344	0.7943E+16	0.1259E+22	0.1585E+06	1.664 ± 0.036	50

this factor would only be necessary in order to compare with other data with different Δt (but caution should be present in this case for the possible fractal nature of the speed signal). Although the speeds are reported in knots, they are converted to m/s (using that 1 knot=0.514 m/s), and then we report the PDI in m^3/s^2 .

Reference [23] studied the statistics of the PDI (defined for individual events, in contrast to Emanuel's work) in 4 different ocean basins for several time periods. The results showed a rapid, perhaps exponential, decay at the tail, but a body of the distribution compatible with a power law, for 1 or 2 orders of magnitude, with exponents close to one. The connection with SOC phenomena was discussed in Ref. [33]. The method used was again a variation of the Clauset's one, introducing an upper cutoff and additional restrictions to the variations of the parameters. Here we revisit this problem, trying to use updated data (whenever it has been possible), and applying the method which is the subject of this paper to $x = PDI$.

The data has been downloaded from the National Hurricane Center (NHC) of NOAA, for the North Atlantic and the Northeastern Pacific [82, 83] and from the Joint Typhoon Warning Center (JTWC) of the US Navy [84, 85] for the Northwestern Pacific, the Southern Hemisphere (com-




FIG. 6: Estimated probability densities and corresponding power-law fits of the seismic moment M of shallow earthquakes in the worldwide CMT catalog and of the estimated M in the Southern California catalog.

prising the Southern Indian and the Southwestern Pacific), and the Northern Indian Ocean. The abbreviation, time span, and number of events for each basin are: NATl, 1966–2011, 532; EPac, 1966–2011, 728; WPac, 1986–2011, 690; SHem, 1986–2007 (up to May), 523; NInd, 1986–2007, 110. The latter case was not studied in any of the previous works.

The results for a truncated power law maximizing N , shown in Table IV and Fig. 7, are in agreement with those of Ref. [23]. In general, exponents are close but above 1, except for the Northwestern Pacific, where $\alpha \simeq 0.96$, and for the North Indian, where α is substantially higher than one. We consider that this method performs rather well. It would be interesting to test if universality can nevertheless hold (the high value for the North Indian Ocean is based in much less data than for the rest of basins), or if there is some systematic bias in the value of the exponents (the protocols of the NHC and the JTWC are different, and the satellite coverage of each basin is

also different).




FIG. 7: Estimated probability densities of the PDI of tropical cyclones in 5 ocean basins, together with their power-law fits. The values of the densities are multiplied by 1, 10^2 , 10^4 , 10^6 , and 10^8 , for clarity sake.

If a non-truncated power law is fit to the data, the fits turn out to be rather short, with a high exponent (up to 6) describing the tail of the distribution (except for the Southern Hemisphere, where no such tail is apparent). We do not give any relevance to these results, as other alternatives, as

TABLE IV: Results of the truncated power-law fits, maximizing N , for the PDI in the 5 ocean basins with tropical-cyclone activity, for different values of p_c .

basin	N	a (m^3/s^2)	b (m^3/s^2)	b/a	$\alpha \pm \sigma$	p_c
EPac	637	0.1259E+10	0.7943E+11	63	1.094 ± 0.033	0.10
NAtl	417	0.1995E+10	0.7943E+11	40	1.168 ± 0.047	0.10
SHem	523	0.1259E+10	0.1259E+12	100	1.108 ± 0.034	0.10
WPac	637	0.5012E+09	0.1259E+12	251	0.957 ± 0.025	0.10
NInd	102	0.7943E+09	0.1995E+12	251	1.520 ± 0.077	0.10
EPac	571	0.1995E+10	0.7943E+11	40	1.149 ± 0.039	0.20
NAtl	417	0.1995E+10	0.7943E+11	40	1.168 ± 0.047	0.20
SHem	523	0.1259E+10	0.1259E+12	100	1.108 ± 0.033	0.20
WPac	637	0.5012E+09	0.1259E+12	251	0.957 ± 0.025	0.20
NInd	102	0.7943E+09	0.1259E+12	158	1.490 ± 0.077	0.20
EPac	571	0.1995E+10	0.7943E+11	40	1.149 ± 0.040	0.50
NAtl	417	0.1995E+10	0.7943E+11	40	1.168 ± 0.045	0.50
SHem	465	0.1995E+10	0.1259E+12	63	1.132 ± 0.040	0.50
WPac	637	0.5012E+09	0.1259E+12	251	0.957 ± 0.024	0.50
NInd	86	0.7943E+09	0.1259E+11	16	1.323 ± 0.139	0.50

for instance a simple exponential tail, have to be considered [80, 86]. Coming back to a truncated power law, but maximizing the log-range, the algorithm sometimes fits the power law in the body of the distribution (with exponent close to 1) and for some other times the algorithm goes to the fast-decaying tail. So the method of maximizing b/a is not useful for this data.

D. Area of forest fires

The statistics of the size of forest fires was an intense topic of research since the introduction of the concept of SOC, at the end of the 1980's, but only from the point of view of cellular-automaton models. Real data analysis had to wait several years [20, 87], leading to power-law distributions, more or less in agreement with the models. Here we are particularly interested in a dataset from Italy, for which a power-law distribution of sizes was ruled out [88]. Instead, a lognormal tail was

TABLE V: Results of the fits for the burned area of the $N_{tot} = 36\,748$ fires recorded in the Italian catalog, for different p_c . The cases of a non-truncated power law and a truncated power law, maximizing N , are shown.

N	a (ha)	b (ha)	b/a	$\alpha \pm \sigma$	p_c
51	794	∞	∞	2.880 ± 0.276	0.10
51	794	∞	∞	2.880 ± 0.277	0.20
51	794	∞	∞	2.880 ± 0.277	0.50
168	316	7943	25	2.353 ± 0.117	0.10
148	316	1259	4	2.075 ± 0.213	0.20
51	794	79430	100	2.870 ± 0.281	0.50

proposed for the fire size probability density.

The data considered in Ref. [88], and reanalyzed in this study, comes from the Archivio Incendi Boschivi (AIB) fire catalog compiled by the Italian Corpo Forestale dello Stato [89]. The subcatalog to which we restrict covers all Italy and spans the 5-year period 1998-2002, containing 36 748 fires. The size of each fire is measured by the burned area A , in hectares, with $1 \text{ ha} = 10^4 \text{ m}^2$. In this subsection we analyze the case of $x = A$.

The results in Table V and Fig. 8 show that a pure (non-truncated) power law is only acceptable (in the sense of non-rejectable) for the rightmost part of the tail of the distribution, comprising less than one order of magnitude. It is very indicative that only 51 data are in the possible power-law tail. Therefore, we disregard this power-law behavior as spurious and expect that other distributions can yield a much better fit (not in order of the quality of the fit but regarding the number of data it spans). This seems in agreement with other analyses of forest-fire data [6, 54]. If a truncated power-law is considered, fitted by maximizing the number of data, the results are not clearly better, as seen in the table. Moreover, there is considerable variation with the value of p_c . So, we do not give any relevance to such power-law fits. Finally, the method of maximizing b/a yields the same results as for the non-truncated power law (except by the fact that the exponents are slightly smaller, not shown). In order to provide some evidence for the better adequacy of the lognormal tail in front of the power-law tail for these data, it would be interesting to apply an adaptation of the test explained in Ref. [90].




FIG. 8: Estimated probability density of the area of fires in the Italian catalog, together with the power-law fits. In contrast to the previous datasets analyzed, we consider these power-law fits as irrelevant.

E. Waiting time between earthquakes

The temporal properties of earthquakes has been a subject relatively little studied (at least in comparison with the size of earthquakes). It is true that the Omori law has been known for more than 100 years [91, 92], and that this is a law extremely important in order to assess the hazard of aftershocks after a big event, but the Omori law looks at time properties in a very coarse-grained way, as it only provides the number of events in relatively large time windows. Thus, no information on the fine time structure of seismicity is provided, at least directly.

The situation has changed in the last decade, since the seminal study of Bak et al. [40], who found a unified scaling law for earthquake waiting-time distributions. They took Southern California and divided it in different areas, and computed the time between consecutive earthquakes for each area. So, if t_i^j denotes the time of occurrence of the i -th earthquake in area j , the corresponding waiting time τ_i^j is

$$\tau_i^j = t_i^j - t_{i-1}^j.$$

The key point is that all the resulting waiting times were added to the same distribution (and not

to a different distribution j for each area). Subsequently, the unified scaling law was found to be valid for other regions of the world [42]. The shape of the resulting probability density corresponds to a double power law, one for small waiting times, associated to a mixture of scales of waiting times due to the Omori law, and another for large waiting times, due to spatial heterogeneity arising from the mixture of different areas with different seismic rates [40–42, 93]. The first exponent was found to be close to 1, whereas the second one was about 2.2; the fits were done by means of the nonlinear least-squares Marquardt-Levenberg algorithm from `gnuplot`, applied to the logarithm of the log-binned empirical density. Here we apply our more sophisticated method to updated data for Southern California seismicity, with $x = \tau$.

We use again the relocated Southern California catalog of Shearer et al. [77, 78], but starting in 1984 and ending in June 30th, 2011. This is to avoid some holes in the catalog for preceding years. As for earthquake sizes, the occurrence takes place in a rectangle of coordinates ($122^\circ\text{W}, 30^\circ\text{N}$), ($113^\circ\text{W}, 37.5^\circ\text{N}$). This rectangle is divided into equal parts both in the West-East axis and in the South-North axis, in such a way that we consider a number of subdivisions of 4×4 , 8×8 , 16×16 , and 32×32 . The waiting times for events of magnitude $m \geq 2$ in each of these subdivisions are computed as explained above, resulting in a number of data between 103 000 and 104 000 in all cases.

For a non-truncated power law, the results are only coherent with the previous reported ones (exponent around 2.2) for the intermediate cases, i.e., 8×8 and 16×16 , see Table VI and Fig. 9. The disagreement for the other cases can easily be explained. For 4×4 , the number of resulting subdivisions, 16, seems rather small. As mentioned, in Ref. [93] the power-law tail was explained in terms of a power-law mixture of exponentials; so, with only 16 regions is possible that the asymptotic behavior is still not reached. On the other hand, the effect of the finite duration of the catalog is visible in the 32×32 data. Due to the scaling behavior of the distributions [41, 42], the possible power-law tail in this case is displaced to larger waiting times; but the time span of the catalog, about 10^{10} s, clearly alters this power law, which starts to bend at about 10^9 s. Thus, we conclude that a power-law exponent of about $\alpha \simeq 2.2$ or 2.3 indeed exists, provided that the number of spatial subdivisions is high enough and the temporal extension of the catalog is large enough.

When a truncated power-law is fitted, using the method of maximizing the number of data N , the other power law emerges, but for a range shorter than what the plot of the densities suggests. The exponent is in a range from 0.95 to 1.04 (except for the 4×4 cases, in which it is a bit

smaller). The largest log-range is 100, i.e., two decades. The graphical representation of the density seems to indicate that the possible power law is influenced by the effect of two crossovers, one for large waiting times, associated to a change in exponent, and another one for smaller times, where the distribution becomes flat. Finally, the method of fitting which maximizes the log-range leads to results that are similar to the non-truncated power-law case, although sometimes intervals corresponding to very small times are selected. The latter results have no physical meaning, as correspond to times below 1 s, i.e., below the error in the determination of the occurrence time.




FIG. 9: Estimated probability densities and corresponding power-law fits for the waiting times of $m \geq 2$ in the Southern California catalog, for different spatial subdivisions. The values of the density are multiplied by factors 1, 10, 100, and 1000, for clarity sake.

TABLE VI: Results of the fits with a non-truncated power law and a truncated power law, maximizing N , for earthquake waiting times calculated for different subdivisions of Southern California. Different minimum p -values are shown. The total number of data is above 103 000 in any case.

Subdivisions	N	a (s)	b (s)	b/a	$\alpha \pm \sigma$	p_c
16×16	542	0.3162E+08	∞	∞	2.324 ± 0.056	0.10
32×32	67	0.3162E+09	∞	∞	4.404 ± 0.405	0.10
4×4	124	0.5012E+07	∞	∞	1.921 ± 0.085	0.10
8×8	1671	0.3162E+07	∞	∞	2.198 ± 0.031	0.10
16×16	542	0.3162E+08	∞	∞	2.324 ± 0.056	0.20
32×32	67	0.3162E+09	∞	∞	4.404 ± 0.403	0.20
4×4	124	0.5012E+07	∞	∞	1.921 ± 0.085	0.20
8×8	1671	0.3162E+07	∞	∞	2.198 ± 0.031	0.20
16×16	24	0.3162E+09	∞	∞	4.106 ± 0.703	0.50
32×32	67	0.3162E+09	∞	∞	4.404 ± 0.449	0.50
4×4	77	0.7943E+07	∞	∞	1.856 ± 0.098	0.50
8×8	322	0.1259E+08	∞	∞	2.231 ± 0.070	0.50
16×16	44178	7943	0.7943E+06	100	0.956 ± 0.004	10
32×32	43512	1259	0.1995E+06	158	1.029 ± 0.003	10
4×4	38765	1995	0.5012E+05	25	0.867 ± 0.006	10
8×8	39851	316	0.1995E+05	63	0.987 ± 0.004	10
16×16	39481	7943	0.5012E+06	63	0.950 ± 0.005	20
32×32	39654	1259	0.1259E+06	100	1.033 ± 0.004	20
4×4	38765	1995	0.5012E+05	25	0.867 ± 0.006	20
8×8	39851	316	0.1995E+05	63	0.987 ± 0.004	20
16×16	39481	7943	0.5012E+06	63	0.950 ± 0.005	50
32×32	39654	1259	0.1259E+06	100	1.033 ± 0.004	50
4×4	34113	3162	0.5012E+05	16	0.864 ± 0.007	50
8×8	39851	316	0.1995E+05	63	0.987 ± 0.004	50

V. CONCLUSIONS

For power-law distributions, the fitting and testing the goodness-of-fit is a difficult but very relevant problem in complex-systems science, in general, and in geoscience in particular. The most critical step is to select, automatically (without introducing any subjective bias), where the power-law regime starts and where it ends. We have explained in detail a conceptually simple but somewhat laborious procedure in order to overcome some difficulties previously found in the method introduced by Clauset et al. [54, 56]. Our method is summarized in fitting by maximum likelihood estimation and testing the goodness of fit by the Kolmogorov-Smirnov statistic, using Monte Carlo simulations. Although these steps are in common with the Clauset et al.'s recipe, the key difference is in the criterion of selection of the power-law range. Despite the many steps of these procedures, ours can be easily implemented, and the resulting algorithms run very fast in current computers. We also have explained how to estimate properly the probability density of a random variable which has a power law or a fat-tail distribution. This is important to draw graphical representations of the results of the fitting (specially in Fig. 5) but it is not required to perform neither the fits nor the goodness-of-fit tests.

The performance of the method is quite good, as checked in synthetic power-law datasets, and the results of the analysis of previously reported power laws are very consistent. We confirm a very broad power-law tail in the distribution of the half-lives of the radionuclides, with exponent $\alpha = 1.09$, as well as other power-law regimes in the body of the distribution. The results for the power-law exponent of the distribution of seismic moments worldwide and in Southern California are in agreement with previous estimates, but in addition our method provides a reliable way to determining the minimum seismic-moment value for which the Gutenberg-Richter law holds. This can be useful to check systematically for the completeness thresholds of seismic catalogs. For the energy dissipated by tropical cyclones, measured roughly through the PDI, we confirm the power-law behavior in the body of the distribution previously reported, with exponents close to one. We also survey new results for the Southern Indian Ocean, but with a higher power-law exponent. In contrast, for the case of the area affected by forest fires in an Italian catalog, we obtain power-law-distributed behavior only for rather small windows of the burnt area, containing few number of data. Finally, for the waiting times between earthquakes in different subdivisions of Southern California we conclude that the power-law behavior of the tail is very delicate, affected either by a small number of subdivisions, when the size of those is large, or by the finite duration of the

record, which introduces a sharp decay of the distribution when their number is high. For the body of the distribution another power law is found, but the range is limited by crossovers below and above it. Naturally, our method can be directly applied to the overwhelming number of fat-tailed distributions reported during the last decades in geoscience.

Acknowledgements

Our development of the fitting and goodness-of-fit testing method presented here was originated from a collaboration with O. Peters, and was influenced by comments from L. A. N. Amaral and R. D. Malmgren. Research in hurricanes and half-lives started jointly with A. Ossó and F. Font, respectively. We got the references for the law of large numbers from J. Ll. Solé. L. Telesca and R. Lasaponara provided permission to use their fire catalog. We are particularly indebted to Isabel Serra, from whom we have learnt a lot, and to Per Bak, who is the ultimate ancestor of all this. Research projects providing some money to us were FIS2009-09508, from the disappeared MICINN, and 2009SGR-164, from the Generalitat de Catalunya.

-
- [1] N. L. Johnson, S. Kotz, and N. Balakrishnan. *Continuous Univariate Distributions*, volume 1. Wiley-Interscience, New Jersey, 2nd edition, 1994.
 - [2] M. Evans, N. Hastings, and B. Peacock. *Statistical Distributions*. Wiley-Interscience, New York, 3rd edition, 2000.
 - [3] N. L. Johnson, A. W. Kemp, and S. Kotz. *Univariate Discrete Distributions*. Wiley-Interscience, New Jersey, 3rd edition, 2005.
 - [4] K. Christensen and N. R. Moloney. *Complexity and Criticality*. Imperial College Press, London, 2005.
 - [5] A. Corral. Scaling and universality in the dynamics of seismic occurrence and beyond. In A. Carpinteri and G. Lacidogna, editors, *Acoustic Emission and Critical Phenomena*, pages 225–244. Taylor and Francis, London, 2008.
 - [6] M. E. J. Newman. Power laws, Pareto distributions and Zipf’s law. *Cont. Phys.*, 46:323–351, 2005.
 - [7] H. Takayasu. *Fractals in the Physical Sciences*. Manchester University Press, Manchester, 1989.
 - [8] A. N. Kolmogorov. *Foundations of the theory of probability*. Chelsea Pub. Co., New York, 2nd edition, 1956.

- [9] A. N. Shiryaev. *Probability*. Springer-Verlag, New York, 2nd edition, 1996.
- [10] J.-P. Bouchaud and A. Georges. Anomalous diffusion in disordered media: statistical mechanisms, models and physical applications. *Phys. Rep.*, 195:127–293, 1990.
- [11] I. B. Aban, M. M. Meerschaert, and A. K. Panorska. Parameter estimation for the truncated Pareto distribution. *J. Am. Stat. Assoc.*, 101:270–277, 2006.
- [12] S. M. Burroughs and S. F. Tebbens. Upper-truncated power laws in natural systems. *Pure Appl. Geophys.*, 158:741–757, 2001.
- [13] Y. Y. Kagan. Seismic moment distribution revisited: I. statistical results. *Geophys. J. Int.*, 148:520–541, 2002.
- [14] H. Kanamori and E. E. Brodsky. The physics of earthquakes. *Rep. Prog. Phys.*, 67:1429–1496, 2004.
- [15] T. Utsu. Representation and analysis of earthquake size distribution: a historical review and some new approaches. *Pure Appl. Geophys.*, 155:509–535, 1999.
- [16] A. Corral and F. Font-Clos. Criticality and self-organization in branching processes: application to natural hazards. In M. Aschwanden, editor, *Self-Organized Critical Phenomena*, 2013.
- [17] B. D. Malamud. Tails of natural hazards. *Phys. World*, 17 (8):31–35, 2004.
- [18] F. Lahaie and J. R. Grasso. A fluid-rock interaction cellular automaton of volcano mechanics: Application to the Piton de la Fournaise. *J. Geophys. Res.*, 103 B:9637–9650, 1998.
- [19] L. McClelland, T. Simkin, M. Summers, E. Nielsen, and T. C. Stein, editors. *Global volcanism 1975-1985*. Prentice Hall, Englewood Cliffs, 1989.
- [20] B. D. Malamud, J. D. A. Millington, and G. L. W. Perry. Characterizing wildfire regimes in the United States. *Proc. Natl. Acad. Sci. USA*, 102:4694–4699, 2005.
- [21] O. Peters, A. Deluca, A. Corral, J. D. Neelin, and C. E. Holloway. Universality of rain event size distributions. *J. Stat. Mech.*, P11030, 2010.
- [22] O. Peters, C. Hertlein, and K. Christensen. A complexity view of rainfall. *Phys. Rev. Lett.*, 88:018701, 2002.
- [23] A. Corral, A. Ossó, and J. E. Llebot. Scaling of tropical-cyclone dissipation. *Nature Phys.*, 6:693–696, 2010.
- [24] M. P. Freeman and N. W. Watkins. The heavens in a pile of sand. *Science*, 298:979–980, 2002.
- [25] S. M. Burroughs and S. F. Tebbens. Power-law scaling and probabilistic forecasting of tsunami runup heights. *Pure Appl. Geophys.*, 162:331–342, 2005.
- [26] J. P. Sethna, K. A. Dahmen, and C. R. Myers. Crackling noise. *Nature*, 418:242–250, 2001.

- [27] M. Aschwanden. SOC systems in astrophysics. In M. Aschwanden, editor, *Self-Organized Critical Phenomena*, 2013.
- [28] E. P. White, B. J. Enquist, and J. L. Green. On estimating the exponent of power-law frequency distributions. *Ecol.*, 89:905–912, 2008.
- [29] P. Bak. *How Nature Works: The Science of Self-Organized Criticality*. Copernicus, New York, 1996.
- [30] H. J. Jensen. *Self-Organized Criticality*. Cambridge University Press, Cambridge, 1998.
- [31] G. Pruessner. *Self-Organised Criticality: Theory, Models and Characterisation*. Cambridge University Press, Cambridge, 2012.
- [32] D. Sornette. *Critical Phenomena in Natural Sciences*. Springer, Berlin, 2nd edition, 2004.
- [33] A. Corral. Tropical cyclones as a critical phenomenon. In J. B. Elsner, R. E. Hodges, J. C. Malmstadt, and K. N. Scheitlin, editors, *Hurricanes and Climate Change: Volume 2*, pages 81–99. Springer, Heidelberg, 2010.
- [34] O. Peters and K. Christensen. Rain viewed as relaxational events. *J. Hidrol.*, 328:46–55, 2006.
- [35] O. Peters and J. D. Neelin. Critical phenomena in atmospheric precipitation. *Nature Phys.*, 2:393–396, 2006.
- [36] R. Dickman. Rain, power laws, and advection. *Phys. Rev. Lett.*, 90:108701, 2003.
- [37] M. Mitzenmacher. A brief history of generative models for power law and lognormal distributions. *Internet Math.*, 1 (2):226–251, 2004.
- [38] M. Baiesi, M. Paczuski, and A. L. Stella. Intensity thresholds and the statistics of the temporal occurrence of solar flares. *Phys. Rev. Lett.*, 96:051103, 2006.
- [39] G. Boffetta, V. Carbone, P. Giuliani, P. Veltri, and A. Vulpiani. Power laws in solar flares: Self-organized criticality or turbulence? *Phys. Rev. Lett.*, 83:4662–4665, 1999.
- [40] P. Bak, K. Christensen, L. Danon, and T. Scanlon. Unified scaling law for earthquakes. *Phys. Rev. Lett.*, 88:178501, 2002.
- [41] A. Corral. Local distributions and rate fluctuations in a unified scaling law for earthquakes. *Phys. Rev. E*, 68:035102, 2003.
- [42] A. Corral. Universal local versus unified global scaling laws in the statistics of seismicity. *Physica A*, 340:590–597, 2004.
- [43] J. A. Wanliss and J. M. Weygand. Power law burst lifetime distribution of the SYM-H index. *Geophys. Res. Lett.*, 34:L04107, 2007.
- [44] A. Corral. Long-term clustering, scaling, and universality in the temporal occurrence of earthquakes.

- Phys. Rev. Lett.*, 92:108501, 2004.
- [45] A. Corral. Statistical tests for scaling in the inter-event times of earthquakes in California. *Int. J. Mod. Phys. B*, 23:5570–5582, 2009.
 - [46] E. L. Geist and T. Parsons. Distribution of tsunami interevent times. *Geophys. Res. Lett.*, 35:L02612, 2008.
 - [47] A. Saichev and D. Sornette. “Universal” distribution of interearthquake times explained. *Phys. Rev. Lett.*, 97:078501, 2006.
 - [48] A. Corral. Comment on “Do earthquakes exhibit self-organized criticality?”. *Phys. Rev. Lett.*, 95:159801, 2005.
 - [49] M. Paczuski, S. Boettcher, and M. Baiesi. Interoccurrence times in the Bak-Tang-Wiesenfeld sandpile model: A comparison with the observed statistics of solar flares. *Phys. Rev. Lett.*, 95:181102, 2005.
 - [50] A. Corral. Universal earthquake-occurrence jumps, correlations with time, and anomalous diffusion. *Phys. Rev. Lett.*, 97:178501, 2006.
 - [51] J. Davidsen and M. Paczuski. Analysis of the spatial distribution between successive earthquakes. *Phys. Rev. Lett.*, 94:048501, 2005.
 - [52] K. R. Felzer and E. E. Brodsky. Decay of aftershock density with distance indicates triggering by dynamic stress. *Nature*, 441:735–738, 2006.
 - [53] H. Bauke. Parameter estimation for power-law distributions by maximum likelihood methods. *Eur. Phys. J. B*, 58:167–173, 2007.
 - [54] A. Clauset, C. R. Shalizi, and M. E. J. Newman. Power-law distributions in empirical data. *SIAM Rev.*, 51:661–703, 2009.
 - [55] M. L. Goldstein, S. A. Morris, and G. G. Yen. Problems with fitting to the power-law distribution. *Eur. Phys. J. B*, 41:255–258, 2004.
 - [56] A. Corral, F. Font, and J. Camacho. Non-characteristic half-lives in radioactive decay. *Phys. Rev. E*, 83:066103, 2011.
 - [57] A. Corral, G. Boleda, and R. Ferrer-i-Cancho. *in preparation*, 2012.
 - [58] A. Corral, A. Deluca, and R. Ferrer i Cancho. A practical recipe to fit discrete power-law distributions. *ArXiv*, 1209:1270, 2012.
 - [59] S. M. Ross. *A First Course in Probability*. Prentice Hall, Englewood Cliffs, 2002.
 - [60] S. Pueyo and R. Jovani. Comment on “a keystone mutualism drives pattern in a power function”. *Science*, 313:1739c–1740c, 2006.

- [61] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes in FORTRAN*. Cambridge University Press, Cambridge, 2nd edition, 1992.
- [62] G. Casella and R. L. Berger. *Statistical Inference*. Duxbury, Pacific Grove CA, 2nd edition, 2002.
- [63] R. Durrett. *Probability: Theory and Examples*. Cambridge University Press, Cambridge, 4th edition, 2010.
- [64] J. D. Kalbfleisch and R. L. Prentice. *The Statistical Analysis of Failure Time Data*. Wiley, Hoboken, NJ, 2nd edition, 2002.
- [65] J. Baró and E. Vives. Analysis of power-law exponents by maximum-likelihood maps. *Phys. Rev. E*, 85:066121, 2012.
- [66] L. Devroye. *Non-Uniform Random Variate Generation*. Springer-Verlag, New York, 1986.
- [67] R. D. Malmgren, D. B. Stouffer, A. E. Motter, and L. A. N. Amaral. A Poissonian explanation for heavy tails in e-mail communication. *Proc. Natl. Acad. Sci. USA*, 105:18153–18158, 2008.
- [68] M. Boguñá and A. Corral. Long-tailed trapping times and Lévy flights in a self-organized critical granular system. *Phys. Rev. Lett.*, 78:4950–4953, 1997.
- [69] A. Corral. Point-occurrence self-similarity in crackling-noise systems and in other complex systems. *J. Stat. Mech.*, P01022, 2009.
- [70] J. Klafter, M. F. Shlesinger, and G. Zumofen. Beyond Brownian motion. *Phys. Today*, 49(2):33–39, 1996.
- [71] S. Hergarten. *Self-Organized Criticality in Earth Systems*. Springer, Berlin, 2002.
- [72] B. W. Silverman. *Density Estimation for Statistics and Data Analysis*. Chapman and Hall, New York, 1986.
- [73] K. S. Krane. *Introductory Nuclear Physics*. Wiley, New York, 1988.
- [74] S. Y. F. Chu, L. P. Ekström, and R. B. Firestone. Lund/LBNL Nuclear Data Search webpage. <http://nucleardata.nuclear.lu.se/nucleardata/toi>, version 2.0, February 1999.
- [75] B. Gutenberg and C. F. Richter. Frequency of earthquakes in California. *Bull. Seismol. Soc. Am.*, 34:185–188, 1944.
- [76] Y. Ben-Zion. Collective behavior of earthquakes and faults: continuum-discrete transitions, progressive evolutionary changes, and different dynamic regimes. *Rev. Geophys.*, 46:RG4006, 2008.
- [77] P. Shearer, E. Hauksson, and G. Lin. Southern California hypocenter relocation with waveform cross-correlation, part 2: Results using source-specific station terms and cluster analysis. *Bull. Seismol. Soc. Am.*, 95:904–915, 2005. <http://www.data.scec.org/ftp/catalogs/SHLK/>, file

SHLK_1_02.txt.

- [78] E. Hauksson, W. Yang, and P. Shearer. Waveform relocated earthquake catalog for southern california. <http://www.data.scec.org/research-tools/alt-2011-dd-hauksson-yang-shearer.html>.
- [79] K. Emanuel. *Divine Wind: the History and Science of Hurricanes*. Oxford University Press, New York, 2005.
- [80] A. Corral and A. Turiel. Variability of North Atlantic hurricanes: seasonal versus individual-event features. In A. S. Sharma, A. Bunde, V. P. Dimri, and D. N. Baker, editors, *Extreme Events and Natural Hazards: the Complexity Perspective*, pages 111–125. Geopress, Washington, 2012.
- [81] K. Emanuel. Increasing destructiveness of tropical cyclones over the past 30 years. *Nature*, 436:686–688, 2005.
- [82] B. R. Jarvinen, C. J. Neumann, and M. A. S. David. A tropical cyclone data tape for the North Atlantic basin, 1886-1983: contents, limitations, and uses. <http://www.nhc.noaa.gov/pdf/NWS-NHC-1988-22.pdf>, 1988.
- [83] National Hurricane Center. <http://www.nhc.noaa.gov/pastall.shtml#hurdat>.
- [84] J.-H. Chu, C. R. Sampson, A. S. Levine, and E. Fukada. The Joint Typhoon Warning Center tropical cyclone best-tracks, 1945-2000. https://metocph.nmci.navy.mil/jtwc/best_tracks/TC_bt_report.html, 2002.
- [85] JTWC Annual Tropical Cyclone Reports. http://www.usno.navy.mil/NOOC/nmfc-ph/RSS/jtwc/best_tracks.
- [86] J. del Castillo, J. Daoudi, and I Serra. The full-tails gamma distribution applied to model extreme values. *ArXiv*, 1211:0130, 2012.
- [87] B. D. Malamud, G. Morein, and D. L. Turcotte. Forest fires: An example of self-organized critical behavior. *Science*, 281:1840–1842, 1998.
- [88] A. Corral, L. Telesca, and R. Lasaponara. Scaling and correlations in the dynamics of forest-fire occurrence. *Phys. Rev. E*, 77:016101, 2008.
- [89] Corpo Forestale dello Stato. <http://www.corpoforestale.it>.
- [90] J. del Castillo and P. Puig. The best test of exponentiality against singly truncated normal alternatives. *J. Am. Stat. Assoc.*, 94:529–532, 1999.
- [91] T. Utsu. Statistical features of seismicity. In W. H. K. Lee, H. Kanamori, P. C. Jennings, and C. Kisslinger, editors, *International Handbook of Earthquake and Engineering Seismology, Part A*,

pages 719–732. Academic Press, Amsterdam, 2002.

- [92] T. Utsu, Y. Ogata, and R.S. Matsu'ura. The centenary of the Omori formula for a decay law of aftershock activity. *J. Phys. Earth*, 43:1–33, 1995.
- [93] A. Corral and K. Christensen. Comment on “Earthquakes descaled: On waiting time distributions and scaling laws”. *Phys. Rev. Lett.*, 96:109801, 2006.